
Rule WLM358: QUEUING IN IOS WAS A MAJOR CAUSE OF I/O DELAY

Finding: Queuing in the I/O Supervisor (IOSQ) was a major cause of the I/O delay with the volume.

This finding applies only to MVS versions prior to OS/390 Release 3, and to MVS versions with OS/390 Release 3 if I/O Priority Management has **not** been specified.

Impact: This finding may have a MEDIUM IMPACT or HIGH IMPACT on the performance of the device.

Logic flow: The following rules cause this rule to be invoked:

- Rule WLM350: I/O activity may have caused significant delays
- Rule WLM351: I/O activity may have caused significant delays
- Rule WLM352: I/O activity may have caused significant delays for server service class
- Rule WLM353: I/O activity may have caused significant delays for server service class

Discussion: IOSQ time is the time from the issuance of a STARTIO macro until the Start SubChannel (SSCH) instruction is issued. After the STARTIO macro is issued, the software determines whether the device is busy. If the device is not busy with this system, the SSCH instruction is issued. However, if the device is busy with this system, the I/O request is queued. Thus, IOSQ time always means that the device is unable to handle additional requests from this system.

Some small IOSQ time is often unavoidable. However, large IOSQ time imply a situation that should be examined. Large IOSQ times result from (1) too many I/O operations directed to the device or (2) lengthy device response times (perhaps caused by high seeking, high RPS delays, or high PEND time).

The following example illustrates the output from Rule WLM358:

RULE WLM358: DEVICE IOS QUEUING TIME WAS A MAJOR CAUSE OF DASD DELAYS

A major part of the potential I/O delay to the TSO Service Class could be attributed to queuing in the I/O Supervisor (IOSQ). IOSQ time is caused by too many I/O operations directed to the device or lengthy device response times (perhaps caused by high seeking, by high RPS delays, or by high PEND time. Please refer to the WLM Component User Manual for advice on how to minimize device IOSQ time.

Suggestion: Large IOSQ times usually involve the following situations:

- Multiple data sets may be active on the volume. This situation is the most common and easiest to solve. The data sets can be redistributed among different volumes, to eliminate the queuing for the single volume.
- Multiple users may be using the same data set on the volume. Depending upon the data set characteristics, duplicate copies of the data set placed on different volumes may solve the IOSQ problems.
- Multiple application systems may be using the volume experiencing high IOSQ times. In this case, perhaps application redesign or scheduling can solve the problem.
- A particular application (or system function) may be executing I/O to the device faster than the device can respond.
- The overall device response time (PEND, DISC, and CONN) times may be large, such that the device is unable to provide quick response to the I/O requests. This situation will be revealed by large values in the PEND, DISC, or CONN measures.

Depending on the amount of IOSQ time involved, on budget considerations, and on the business importance of the work being delayed, you might consider acquiring Parallel Access Volumes (PAV). The PAV design tends to eliminate IOSQ time.